IN THE SPECIFICATION:

Please amend the Specification as follows:

Please delete paragraph 12 by deleting it in its entirety and replacing it with the following paragraph:

[0012] FIG. 5 is a graph showing the utilization and worst-case transfer rate as a function of the number of nodes in a system;

FIG. 6 is a graph showing the transfer rate of a system with N=4 at various pre-set utilization thresholds;

FIG. 7 is a graph showing the transfer rate when N is varied for the same guaranteed utilization;

FIG. 8 is a graph showing a comparison between the utilization and transfer rates of global and local schemes; and

FIG. 9 is a graph showing the impact of the node capacity on the behavior of the system.

Please amend paragraph 20 as follows:

[0020] FIG. 3 is a flowchart of processing performed when a new physical node is added to system, in accordance with an embodiment of this aspect of the invention. At step 301, the new physical node is added to the network, and, at step 302, a "bootstrap" phase occurs in

which the physical node locates the rest of the peer-to-peer network. For example and without limitation, this can be done by associating a domain name system name with a given distributed hash table network, as described in S. Ratnasamy et al., "A Scalable Content Addressable Network," in Proc. of ACM SIGCOMM, 2001, which is incorporated by reference herein. After locating the peer-to-peer network the new physical node conducts a search for a physical node hosting multiple zones at step 303. This search need not be a global one. Only up to a fixed number of physical nodes need be contacted. If a node hosting at least two zones is found at step 304, then one such zone is transferred, as depicted in step 305, to the new physical node. This reestablishes the invariant that each node has at least one zone. If, however, only physical nodes with one zone are found, then one such zone is selected for what the inventors refer to as an "eager" split. At step 306, the zone is split, generating two children zones, and one of them is transferred to the new physical nodes. To limit the network imbalance, it is preferable to select the largest possible zone for the split in such a case.

Please amend paragraph 22 as follows:

[0022] In accordance with an embodiment of another aspect of the invention, it is advantageous to allow each physical node to hold more slots than would actually fit given its capacity. In other words, the physical node "pretends" that it has more space than is actually available at the physical node. The inventors refer to this technique as "oversubscription." It is observed that if the zones residing on a node are lightly loaded (as most of them are when the system has low utilization), there is enough storage on the node to host extra zones. Consequently, the above-mentioned aggregation technique can be advantageously modified as

follows. In accordance with a preferred embodiment of the invention, each physical node is assigned a number of slots equal to $2 \times N - 1$ slots, where $N = C / SlotSize$ and C is the node's capacity. This is illustrated by FIG. 2. The physical node in FIG. 2 has actual capacity of 300GB where the *SlotSize* is 100GB. Thus the physical node [[210]] ,for example, 110, has three "real" slots 210 and two "virtual" slots 220, each of which are partially filled. Groups of "real" and "virtual" slots are also depicted in FIG. 1 at reference numbers 111, 121, 131, 141.